Mathenautics: Using VR to Visit 3-D Manifolds

Randy Hudson[1]
University of Illinois
Charlie Gunn[2]
Technical University, Berlin
George K. Francis[3], Daniel J. Sandin[4], Thomas A. DeFanti[4]
University of Illinois

Abstract

In most virtual reality applications, 3-d space is a passive, ambient continuum in which the objects of study are placed. When the 3-d space itself is the object of study, as with mathematical manifolds, VR is especially important as a visualization medium. We describe the visualization of such spaces in the CAVE virtual environment.

1 Introduction

Computer graphics has become instrumental in new discoveries in several domains of mathematics. For example, in the study of minimal surfaces computer graphics was indispensable in advancing several proofs and conjectures regarding a new minimal surface [1]. Depending on the emphasis, these new techniques are called either visual or experimental mathematics. Several centers have been founded to further research in this direction, such as the Geometry Center at the University of Minnesota and the SFB-288 lab at Technical University Berlin.

One area of visual mathematics where computer graphics can be helpful is the classification of 3-dimensional manifolds (*3-manifolds*). According to a conjecture of William Thurston (we omit some technical conditions), any 3-manifold can be classified by modeling it on one of eight model geometries (see [4] and related literature). In this paper, we will be concerned with three of these model geometries (the classical cases which exist in every dimension): euclidean (E^3), spherical (S^3) and hyperbolic (H^3). We will use the term manifolds loosely, to include the related spaces known as *orbifolds*, which may contain singular points. For a fuller discussion of the mathematical background and its computer-graphical implementation followed here, see [3]. The next section contains a very abbreviated version of this article.

[1] Academic Information Technologies, Culver Hall, 1025 E. 57th, University of Chicago, rhudson@eecs.uic.edu

[2] Special Research Project 288, "Differential Geometry and Quantum Physics", Math 8-5, Strasse des 17 Juni 136, Technical University, 10623 Berlin, Germany, gunn@sfb288.math.tu-berlin.de

[3] National Center for Supercomputing Applications, University of Illinois at Urbana-Champaign, gfrancis@math.uiuc.edu

[4] University of Illinois at Chicago

2 Previous work on manifold visualization

The visualization of 3-manifolds is not as straightforward as that of 2-manifolds. The 2-torus can be visualized directly as a familiar doughnut-shaped surface by embedding it in a higher (third) dimension. This form of representation can be called the *outsider's* view. However, living within 3-dimensional space, we have no fourth, directly-visible dimension in which to embed or immerse 3-manifolds.

An alternative method for visualizing manifolds which solves this problem is the *insider's* view. This is the view we would see if we were to live inside the manifold. It is constructed as a tessellation of the model geometry by non-overlapping copies of a single tile, or fundamental domain. One copy of this tile represents the underlying manifold; the other copies represent the different ways that light can travel in the manifold to reach the observer's eye. In the case of our two-dimensional torus, the corresponding tessellation covers the euclidean plane with copies of a parallelogram. (To get the outsider's view we take one copy of this parallelogram and roll it up in 3 dimensions to make the torus.) The fundamental domain is replicated via the application of (*a discrete group of*) isometries, in this case, translations in two independent directions. Likewise, the 3-torus is visualized from the inside by applying translations in three independent directions of E^3. (The discrete groups we study can contain other kinds of isometries: glide-reflections, screw motions, rotations and reflections.) Figure 1 shows the insider's view of the 3-torus where the fundamental domain – a cube-like polytope – has been shrunk to improve visibility.

3 Introduction to GeomCAVE

The immersive nature of the insider's view is what makes a virtual environment a better visualization medium for 3-manifolds than the graphics workstation. The Geometry Center has developed an interactive viewer *geomview* based on the graphics library OOGL, capable of visualizing the three model geometries considered here on a workstation. The second author, while at the Geometry Center, developed a related tool, *maniview*, for visualizing the insider's

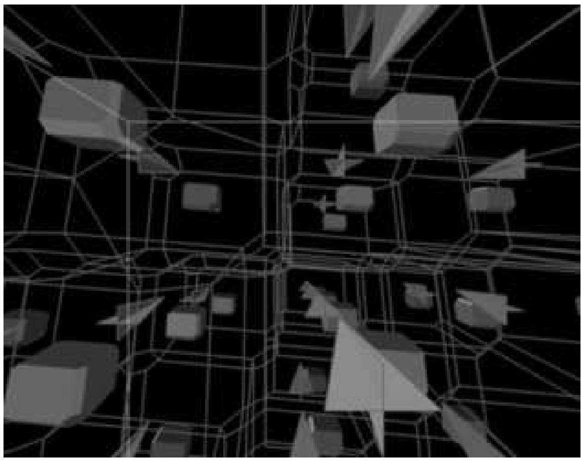

Figure 1: Tessellation of euclidean 3-space by 3-torus.

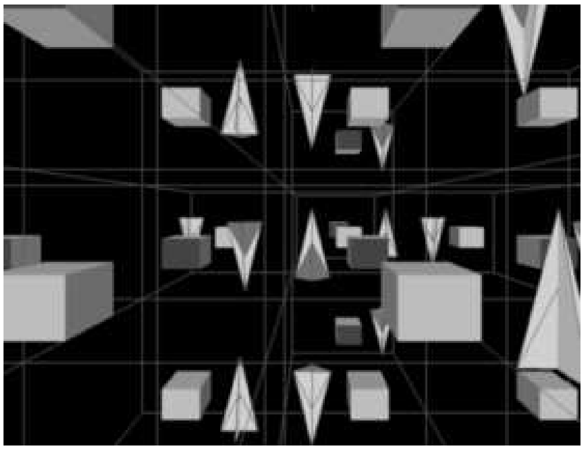

Figure 2: Tessellation of euclidean 3-space by group containing rotation axes.

view of 3-manifolds modeled on these geometries[5]. We have adapted these tools to the CAVE virtual environment developed at the Electronic Visualization Laboratory at the University of Illinois of Chicago [2]. We call this hybrid tool GeomCAVE; using it, the observer can for the first time actually travel within these 3-manifolds and see them from the inside. (After GeomCAVE was developed, more modest alternatives for traveling through hyperbolic and spherical space using the standard CAVE libraries were developed by Ulrike Axen, Glenn Chappell, Chris Hartman, Joanna Mason, Paul McCreary and the fifth author for the Post-Euclidean Walkabout at SIGGRAPH '94. Stuart Levy and Tamara Munzner, from the Geometry Center, have recently expanded this code to read a subset of the OOGL formats directly into the CAVE.)

Modules for the CAVE generally take the form of a single draw routine consisting of GL function calls, which is called regularly from the main CAVE program. Since OOGL maintains its own graphics context, including transform and appearance stacks, which objects rely on when rendering themselves, GeomCAVE had to be slightly more sophisticated. We had to make some CAVE states available to the OOGL context (e.g., which wall is currently being drawn); and we also channeled the navigation data (walking and flying data) through OOGL routines to generate non-euclidean isometries.

The user of GeomCAVE is provided with a menu of icons representing different manifolds; choosing one brings him into the tessellation for that space. The fundamental tile is by default provided by a Dirichlet domain for the underlying group; it is represented once at full scale in wire frame and once at reduced scale as a shaded solid, with corresponding faces a unique color. The reduction in scale was necessary to provide visibility of the whole tessellation, while the coloring provides important information of the structure of the manifold. The observer is represented in the scene by a small dart-shaped object which is also tessellated (figure 2). It points in the direction of the observer's gaze, and its motion and orientation with respect to the fixed geometry of the tessellation provides further structural information. Some of the example spaces contain singular axes; approach to these is signaled by the convergence of multiple copies of the dart to a single point.

The observer can navigate through the space either by physically walking within the CAVE or by flying in the direction of a hand-held wand. He can reset himself to the origin if he gets lost, or can return to the icons to select another space to visit.

[5] geomview and maniview are available via ftp from geom.umn.edu in pub/software

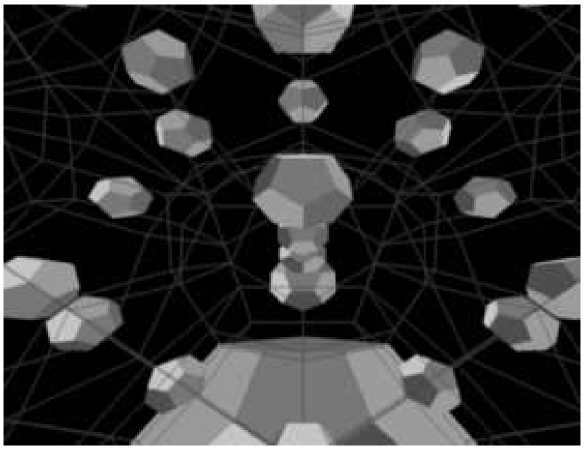

Figure 3: Dodecahedral tessellation of spherical 3-space.

3.1 The example manifolds

Two of the non-euclidean examples are tessellated by regular dodecahedra, a construction impossible in E^3. Because the sum of a triangle's angles in H^3 is less than 180°, a regular dodecahedron with right dihedral angles is possible. H^3 can be tiled with these right, regular dodecahedra in a variety of ways. In GeomCAVE, the observer can verify this experimentally by flying or walking to the common corner of eight dodecahedra and examining these right angles directly. See [3] for an illustration of an H^3 tessellation.

In the same way, it is possible to have a regular dodecahedron in S^3 which has 120-degree dihedral angles (figure 3).

The euclidean 3-manifolds featured in GeomCAVE are both based on the tessellation of euclidean space by a cube, but the discrete groups which perform the tessellation are different. See Figures 1 and 2. It is also possible to explicitly provide geometry to be tessellated instead of using the Dirichlet domain (figure 4).

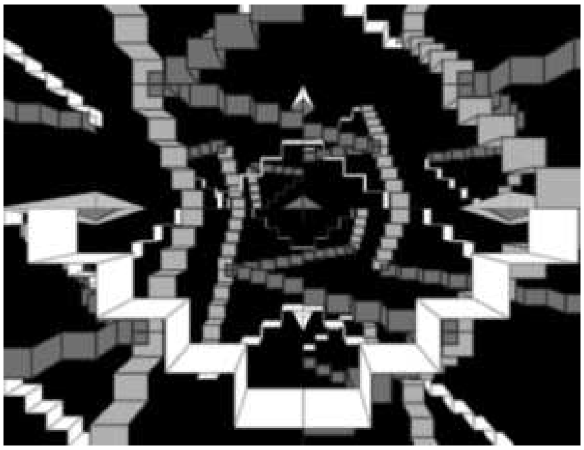

Figure 4: Alternate tessellation of 3-space by the group in Figure 2.

4 GeomCAVE implementation challenges

We encountered several challenges in the implementation of GeomCAVE:

- Conflicting viewing paradigms,
- Non-euclidean navigation in GeomCAVE,
- Manifold navigation in GeomCAVE,
- Stereo, and
- Efficiency measures.

4.1 Conflicting viewing paradigms

OOGL's on-axis perspective projection is applicable in a head-mounted display VR system, where the view planes move with the viewer's eyes, but not in the CAVE, where off-axis projection must be used because the view planes are stationary. This problem applied equally to all geometries. The solution required us to replace the OOGL camera object with an alternative means of transforming from world coordinates to screen coordinates based upon off-axis projection from the observer's position within the CAVE.

4.2 Non-euclidean navigation in GeomCAVE

Euclidean 3-space can be trivially modeled in the euclidean 3-space of the CAVE interior. We discuss the more difficult challenge of mapping hyperbolic geometry into the CAVE; similar remarks apply to the spherical case.

The model of H^3 we use in GeomCAVE the *Beltrami-Klein* (or *projective*) model [3]. In this model, hyperbolic space is modeled as the interior of the unit ball in E^3. Measurements of distance and angle are computed using a different metric; the result is that the unit sphere lies an infinite (hyperbolic) distance from any point within the ball. Hence the unit sphere is called the *sphere at infinity* in this model. Figure 5 shows two views of a tessellation of the two-dimensional hyperbolic plane in this model.

These facts have two consequences for navigation within GeomCAVE. The first follows from the requirement that the observer should not be allowed to leave hyperbolic space by walking (see end of this section). If we want to prevent the navigator from leaving H^3 then it is clear that we must map the CAVE's physical coordinates into the projective model so that it lies entirely within the unit ball. We settled on a scaling value of 0.1, yielding typical corner coordinates of .5, .5, .5, which keeps the CAVE well within the unit sphere. At this value, the dihedral angles between the cave walls, measured hyperbolically, are around 70°. Shrinking the cave yields angles approaching the right angle of Euclidean measurement.

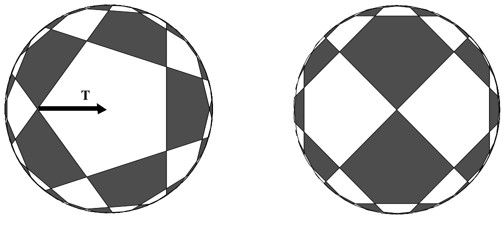

Figure 5: Tessellation of hyperbolic plane by alternately-colored, regular, right-angled pentagons, in the Klein model. Right figure represents a hyperbolic translation of the left by the vector T.

The second consequence affects how navigation is conceptualized. The most natural way to think of flying or walking is that the observer moves through the scene. However, if we follow this model when we implement hyperbolic movement, the result is incorrect. The mistake occurs in the standard construction of the off-axis perspective transformation, which typically contains an implicit euclidean translation to move the observer to the origin. This euclidean translation naturally results in an incorrect image. The solution is to hyperbolically translate the scene past the observer, rather than vice-versa.

In contrast to H^3, which is modeled on the interior of the unit ball, the projective model of S^3 contains all the available points. There is an implicit restriction on the size of the CAVE in model coordinates, since the intrinsic metric of S^3 is finite. The CAVE can not be made larger than a certain size; beyond that scale the size of the cave begins to shrink, just as a circle on the 2-sphere attains a maximum size at the equator and beyond that point gets smaller.

We made the decision to represent H^3 as an inhabitant would experience it, which prohibited the navigator from traveling outside the space. We would like to explore the possibility of euclidean exploration of H^3, so that mathematicians can see how this model "sits" within ordinary space (see Section 6).

4.3 Manifold navigation in GeomCAVE

Though not new with GeomCAVE ([3]), navigation in manifolds is worth mentioning here. One of the challenges unique to manifold exploration involves "staying centered" in the tessellation. Since E^3 and H^3 are infinite (as opposed to S^3, which is a finite space), a complete tessellation would be of infinite extent. Since we can only create a finite tessellation, the possibility exists that the navigator might fly beyond the computed tessellation. The solution adopted here is to "cage" the navigator within the central Dirichlet domain of the group. That is, if in the motion of walking or flying, a wall of the central Dirichlet domain passes by the observer (fixed at the origin), then the observer is moved to an equivalent point lying within the central fundamental domain. That is, the cumulative navigation isometry is multiplied by the group element associated with the crossed face. The resulting isometry maintains the origin within the central Dirichlet domain. With respect to the manifold, this new transform is equivalent to the original isometry, since multiplication by group elements leave the manifold invariant. However, this multiplication may be detectable in our finite implementation: some copies on the edge of the tessellation may appear or disappear. In the ideal implementation (requiring more computer power) these copies are barely visible, either being too small or too foggy.

The alternative, to translate the tessellation to follow the observer, quickly leads in the hyperbolic case to severe numerical problems in the action in the group elements. The result is that the fourth, "homogeneous" coordinate of the transformed vertices grows exponentially large and the dehomogenization operation loses precision. This is avoided by the method outlined above.

4.4 Stereo

Modeling stereo vision presented challenges in the non-euclidean case. We first describe the more familiar solution available in the euclidean setting. The observer and the CAVE have a fixed physical reality which should be mirrored in the models we apply to them. That is, the model coordinates for the navigator are the same as the model coordinates of the CAVE. In particular, the interocular separation of the observer stays at a fixed ratio to the CAVE size. We found empirically that an interocular distance of about 1/100 that of the diagonal of the CAVE is small enough to assure fusion. This translates to a distance of about 2 inches, roughly corresponding to human anatomy. In euclidean space, making the CAVE larger is equivalent to shrinking the scene while keeping the CAVE a constant size. However, in non-euclidean settings, this equivalence no longer holds! In these spaces, there is no change of size without also changing shape. Consequently, it is the CAVE and observer that changes size (and shape!) while the scene remains the same. Of course there is no guarantee of fusion; it may become difficult if the observer in H^3 becomes too large while standing near the fixed geometry; but the danger is no different from the physically observed difficulty of fusing stereo when you move your hand closer to your eyes in everyday life.

The pair of images for the stereo effect is produced by rendering each eye separately as described above by hyperbolically translating the scene to locate the given eye at the origin.

4.5 Efficiency measures

To maintain the frame rate required in VR, we needed to disable the software lighting and shading for non-euclidean scenes (OOGL does lighting in software because of the different metrics of the non-euclidean geometries). We kept the model of the tessellation simple – a wireframe, with simple, solid tiles inside. The discrete group software in OOGL automatically culled the copies of the tessellation which lay outside the viewing frustum of a given wall of the CAVE. Also, we kept the number of layers of the tessellation great enough to produce a sense of depth, but small enough to maintain an adequate frame rate.

5 Evaluation

We have combined the discrete group capabilities of OOGL with VR, the only visualization paradigm for an immersive, direct experience of mathematical spaces, to extend the power of interactive 3-d visualization of such spaces. Access to 3-manifolds via a virtual environment is a significant addition to the tools available for mathematical research and education. For example, as pointed out in section 3, GeomCAVE allows direct observation of interesting properties of non-euclidean spaces, such as the right angles of dodecahedra in hyperbolic space. GeomCAVE immediately makes features of OOGL available in VR, such as a collection of geometric models and discrete group operations. Thus, a mathematician who has built a manifold for viewing in maniview would be able to also explore it in GeomCAVE.

6 Further work

- Implement mixed mode navigation in H^3 (see conclusion of Section 4.2).
- Add more features of maniview:
  - Control over the size and shape of the Dirichlet domain.
  - Control over the depth of the tessellation.
  - As hardware improves, re-activate the software shading and fog effects.
- More sophisticated tools for mathematicians:
  - Connections with existing manifold software (such as *snappea* ([5]).
  - Finer interactive control of the discrete group: selecting subgroups, use of color, deformation of the group.
  - Simulation of dynamical systems in non-euclidean spaces.
  - Extend the coverage to the other five Thurston geometries.
- Experiment with audio tessellation along with the geometric data. The resulting echo patterns could distinguish differently-shaped manifolds.

7 Acknowledgements

We would like to extend special thanks to Stuart Levy of the Geometry Center for his help. Thanks are also due to Mark Phillips and Tamara Munzner, also of the Geometry Center, as well as Louis Kauffman, of the University of Illinois at Chicago.

REFERENCES

[1] Callahan, M.J., Hoffman, D. and Hoffman, J.T. Computer Graphics Tools for the Study of Minimal Surfaces. *Communications of the Association for Computing Machinery 31*, 6 (1988), 648-661.

[2] Cruz-Neira, Carolina, Sandin, Daniel J., DeFanti, Thomas A., Kenyon, Robert V. and Hart, John C. The CAVE: Audio Visual Experience Automatic Virtual Environment. *Communications of the Association for Computing Machinery 35*, 6 (June, 1992), 65-72.

[3] Gunn, Charlie. Discrete Groups and Visualization of Three Dimensional Manifolds. *Computer Graphics 27* (July, 1993), 255-262. Proceedings of SIGGRAPH 1993.

[4] Thurston, William. Three Dimensional Manifolds, Kleinian Groups and Hyperbolic Geometry. *BAMS 19* (1982), 417-431.

[5] Weeks, Jeff. snappea — a MacIntosh application for computing 3-manifolds". (available from ftp@geom.umn.edu).